REMARKS/ARGUMENTS

Reconsideration of this application is respectfully requested.

The rejection of claims 1-8 under 35 U.S.C. §102 as allegedly anticipated by Ramakrishnan '029 is respectfully traversed.

This invention manages the generation of pause frames on a link over which a network device receives data packets from a remote source at the other end of the link. It is known, and is, for example, part of the transmission Standard 802.3 mentioned in the applicant's specification (and in Ramakrishnan), to send on a link, pause frames which have an operating code identifying them as pause frames and a code which indicates the length of time for which the source should cease transmission on the link. Pause frames are generated in order to avoid over-filling of buffer memory, which is necessarily of limited size and limited off-loading rate, and thereby avoid unnecessary complications such as the discard of packets. Although various transmission protocols, such as TCP, enable a system to recover from the loss of packets, in general the technique of 'flow control' using pause frames is a better expedient.

The ordinary way, and illustrated, as will be explained, by Ramakrishnan, of organizing the generation of pause frames is to provide at least one 'threshold' or 'watermark' (the terms are interchangeable) which denotes that a memory approaches a full state. The location of such a watermark is a matter of design choice but usually it is

situated, as indicated by Ramakrishnan, near the top of the level of occupancy, such as 95% of the maximum effective capacity of the memory. The usual scheme is to detect when the occupancy exceeds the threshold and at that point to generate a pause frame (XOFF) with a very long pause time. Usually the pause time is finite (avoiding the need to send a cancelling command) but in network terms very long.

Ramakrishnan generates a pause frame in this manner, as noted in his abstract and the passage on column 8, line 40 to line 54, and most conveniently illustrated in Figure 7. The receiving routine, after the storage of a frame in the receive buffer, stage 704, determines whether the level of occupancy (the buffer level) exceeds the upper threshold. If it does, then a pause frame is generated, stage 710. If the buffer level is less than the upper threshold a further inquiry is made on whether the buffer level is less than the lower threshold. If it is, then a pause frame is generated but as indicated in Ramakrishnan, it is a pause frame of a different type.

As is apparent from Ramakrishnan (e.g., column 7, line 40 to 63), a single pause frame prescribing cessation of transmission by the source is generated if the upper threshold is reached. If the level of occupancy goes below the lower threshold then a frame (XOFF) which has zero pause time is generated.

The problem with Ramakrishnan, as actually discussed in the applicant's specification (e.g., page 2, line 10 to line 23), is that the size of the buffer memory needs

to be rather large. It needs to be able to accommodate the worst case scenario, which arises when the sending rate of the source is high and there is a long round trip propagation delay on the link. It will be understood that the sending source will not actually stop sending frames until it receives the pause frame and so the buffer memory will continue to receive frames after the generation of the 'off' pause frame. In practice this means that the upper watermark needs to allow considerable space for the transmission of these frames so that there is a substantial memory space required 'above' the upper watermark. Further since there is no inhibition of frames until the occupancy level reaches the upper watermark, there are very large amplitude oscillations in the occupancy of the buffer memory.

The present applicant has adopted an approach that is completely different from Ramakrishnan for the generation of pause frames. One basic difference is in the generation of 'off' periods, e.g., pause frames with a code specifying cessation of transmission by the source, when the occupancy is below the relevant watermark, not above it as in Ramakrishnan. Although Ramakrishnan generates a particular type of pause frame when the occupancy is below the lower watermark, these pause frames are not 'off' pause frames; they are pause frames with a zero time and are 'on' frames allowing transmission.

Another fundamental difference is that the applicant's pause frame generator does not only generate a single pause frame. It can generate a repetitive sequence of pause

frames which prescribe alternating periods of transmission and non-transmission from the source. One fundamental advantage of this feature is the ability to limit the rate at which the memory buffer is occupied (e.g., a value determined by the duty cycle of the repetitive sequence of pause frames, that is to say the ratio between the off and on times). Accordingly, the input flow of packets is regulated even when it is permitted.

One basis for the present invention can be seen with reference to applicant's Figure 5. The critical level is not the high watermark but the low watermark. As the level of occupancy of the buffer memory goes below the low watermark, shown at time t4, the repetitive sequence of off and on periods is commenced. It imposes a limitation on the slope of the occupancy curve, in particular between the times t6 and t8. Without the applicant's novel approach, this curve would be very much steeper and the separation between the high watermark and the low watermark would have to be much greater, as indeed it has to be in a system such as described by Ramakrishnan.

Another optical but preferred feature of the applicant's is the choice of when the sequence of off and on periods should end. It may end when the occupancy is at the high watermark, at which time the control system may go into an 'off' mode.

All the claims include distinctive features of the invention and are clearly allowable over the prior art. For example, Ramakrishnan does not disclose the provision of a sequence of pause frames defining alternating on and off periods when the

occupancy of the memory is less than a watermark. All Ramakrishnan does is to provide a single 'off' pause frame if the level of the occupancy is greater than the upper watermark or a single 'on' pause frame if the level of occupancy is below the lower watermark. The only time when Ramakrishnan provides a sequence of frames when the level of occupancy is below a watermark is when he provides a sequence of pause frames with zero pause time. By definition these are not frames which prevent transmission on the link. It is a mis-representation of Ramakrishnan and an act of violence on the language of the claims to read onto Ramakrishnan applicant's generation of an alternating sequence of off and on times when the occupancy is below a watermark. Ramakrishnan only ever generates one frame at a time, as is plainly apparent from his Figure 7 and the corresponding description.

Contrary to the Examiner's assertion beginning on the last line of page 2 of the Office Action, Ramakrishnan does not provide a 'sequence of pause frames wherein said pause frames in the sequence define for a source of said packets alternating periods in which sending of packets on the link is alternately allowed and prevented as claimed.

To repeat, Ramakrishnan provides a single 'off' pulse if the level of occupancy is above the upper threshold. This therefore cannot anticipate claim 1. Ramakrishnan provides a single 'on' pulse if the level of occupancy is below the lower watermark. This also is another reason Ramakrishnan cannot anticipate claim 1. Not even the combination of the two sub-processes is relevant because Ramakrishnan's 'off' pulses are

only generated when the level of occupancy is above the upper watermark. Accordingly, even a distortion of the language and Ramakrishnan, trying to read claim 1 onto an operation in Ramakrishnan where the occupancy of the memory quickly oscillates between the upper and lower thresholds still fails. It should be apparent in any event that the automatic sequence that the applicant provides is nowhere to be found in Ramakrishnan and not even contemplated by him.

With respect to claim 6, basically the same remarks apply. The Examiner's assertions in relation to column 8 and column 9 are clearly erroneous. The Examiner asserts there is 'means responsive to said indication for providing a sequence of pause frames comprising the alternating of x-off and x-on frames. This is not correct. It is not disclosed by Ramakrishnan, who describes the provision of a single x-off frame when the buffer level is above the higher threshold and the generation of a single x-on frame when the occupancy is below the lower level.

With respect to claim 2, the Examiner relies on claims 17 to 19 of Ramakrishnan. However, these claims refer only to a single pause frame not a sequence. Furthermore, it is plainly apparent that the Ramakrishnan x-off frame is generated only when the level of occupancy is 'greater than the upper level threshold'. The Examiner's reliance on column 8, line 47 of Ramakrishnan is misguided. The passages refer to Figure 7, which has already been discussed above. Column 9, lines 17 to 39 refer only to the priority given to a pause frame over other non-pause frames to be sent over the link to the source.

Examiner's reliance on Figure 8 is also misdirected. Figure 8 refers only to the action of the source, i.e., what the source does in response to the pause frames.

The Examiner's objection to applicant's claim 3 is likewise mistaken. Claim 3 refers to the cessation of the sequence of pause frames. As is demonstrated above, Ramakrishnan does not actually generate a sequence as required by parent claim 1. Claim 3 refers to the cessation of the sequence when the occupancy reaches the upper threshold. By contrast, Ramakrishnan commences the sending of a pause frame when his level of occupancy reaches the upper threshold.

With respect to claims 4 and 7, contrary to the assertion of the Examiner, these claims are manifestly not anticipated by Ramkrishnan. What these claims are saying is that the sequence of x-off and x-on frames commences when the occupancy is below the relevant threshold and ceases when the level of occupancy goes above the same threshold. This refers to the (optional) feature that the generation of the sequence of pause frames is stopped a time t7 in Figure 5, that is to say at the low watermark. It is plain that Ramakrishnan's watermarks have to be different and are not the same.

With respect to claims 5 and 8, these claims are directed to the optional feature that when the level of occupancy goes over an upper watermark a pause frame, typically an x-off frame is generated. This feature, by itself, may arguably correspond to Ramakrishnan's generation of the sequence of an x-off frame when the occupancy goes

over the higher watermark. However, this feature is here claimed only in combination with generation of the sequence of x-on and x-off pulses which have been described in parent claims.

The Examiner's attention is also drawn to new claims 9-24. New apparatus claims 9-16 are similar to new method claims 17-24 respectively. New independent claims 9 and 17 can be analogized to original claim 1 in many respects. However, for example, new apparatus claim 9 avoids any "means plus function" recitations and new claim 17 is directed to a method rather than apparatus. Independent claims 12 and 20 specify the termination of the sequence at some selected watermark.

Accordingly, this entire application is now believed to be in allowable form and a formal Notice to that effect is respectfully solicited.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By: 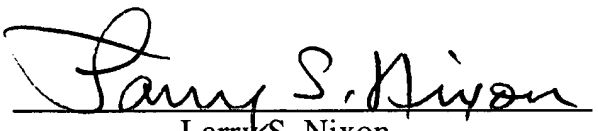

Larry S. Nixon
Reg. No. 25,640

LSN:vc
1100 North Glebe Road, 8th Floor
Arlington, VA 22201-4714
Telephone: (703) 816-4000
Facsimile: (703) 816-4100